ARMY RESEARCH LABORATORY

Developing Effective Human Supervisory Control for Air and Missile Defense Systems

by John K. Hawley and Anna L. Mares

ARL-TR-3742 **February 2006**

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory
Adelphi, MD 20783-1197

ARL-TR-3742 **February 2006**

Developing Effective Human Supervisory Control for Air and Missile Defense Systems

John K. Hawley and Anna L. Mares
Human Research and Engineering Directorate, ARL

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
February 2006	Final	1 Oct 2005 to 30 Sep 2005

4. TITLE AND SUBTITLE
Developing Effective Human Supervisory Control for Air and Missile Defense Systems

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
John K. Hawley and Anna L. Mares

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
U.S. Army Research Laboratory
ATTN: AMSRD-ARL-HR-ME
Ft. Bliss Field Element
Ft. Bliss, TX 79916

8. PERFORMING ORGANIZATION REPORT NUMBER
ARL-TR-3742

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. Army Research Laboratory
Human Research and Engineering directorate
Aberdeen Proving Ground, MD 21005-5425

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

One of the defining properties of the next generation of air and missile defense (AMD) command and control (C2) systems is an increasing reliance on automation. This report is the second in a series of three dealing with human performance and training issues in the development and effective use of automated systems for real-time AMD C2. The first report (Hawley, Mares, & Giammanco, 2005) addresses the impact of automation on air defense operators and the consequences of their role change from traditional operators to supervisory controllers. The present report expands upon that original material and discusses developing effective human supervisory control in AMD C2 systems. Together, these reports are intended as a primer on automation, supervisory control, and effective human performance for commanders, concept developers, system designers, trainers, and other personnel involved with acquisition and use of the next generation of AMD C2 systems.

15. SUBJECT TERMS
Patriot, air and missile defense, command and control, automation, supervisory control

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			John K. Hawley
Unclassified	Unclassified	Unclassified	SAR	46	19b. TELEPHONE NUMBER (*Include area code*) (915) 568-2896

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

List of Tables

INTENTIONALLY LEFT BLANK

Executive Summary

One of the defining properties of current and future air and missile defense (AMD) systems is an increasing reliance on automation. Technology and an increasingly complex operating environment have created a situation where AMD operators must be provided with automated decision support to meet mission objectives. There is a tendency among system developers with little background in human performance issues to assume that automation is innately beneficial. Research in a number of areas suggests, however, that such is not always the case. To begin, automation elevates operators into system monitors rather than active controllers. Operators are thus removed from moment-to-moment, active control and become monitors and managers of subordinate automated processes. It is a well-established fact that humans make very poor system monitors.

Beyond classical vigilance, there are other problems associated with the role change from traditional to supervisory control. Research and operational experience indicate, for example, that automation does not replace human operator tasks. Rather, automation changes the nature of the work that operators do. And it does this in ways that are often unanticipated by system developers and users. Other problems associated with supervisory control generally fall into one of two categories: (1) loss of situation awareness (SA) and (2) skill impairment. SA is important because it has been shown to be a key determinant of decision quality in battle command. Automation in and of itself does not prevent operators from establishing and maintaining SA or contribute to skill impairment. However, improper implementation coupled with inadequate or inappropriate training can make it more difficult for operators to establish and maintain SA and contribute to skill development and retention problems. The preponderance of theory and empirical evidence suggests that the job of supervisory controller is quite different from that of a traditional operator. To maintain system effectiveness, these differences must be reflected in system design, performance support features (i.e., job aids), and operator training.

Much of contemporary real-time automation applied to C2 illustrates what has been referred to as the Catch-22 of human supervisory control: Automation has been introduced because it can do the job better than human controllers, but humans have been left in the control loop to "monitor" that the automated system is performing correctly and override the automation when it is "wrong." The unstated assumption is that operators can properly decide when the automation's decisions should be overridden. Humans are expected to compensate for machine unreliability, but they suffer from a variety of cognitive limitations and vulnerabilities that make it nearly impossible to meet this expectation. A number of automation researchers have thus concluded that while the risks associated with automation unreliability can never be eliminated entirely, they can be managed more effectively through a number of positive actions directed at

supporting and enhancing effective human supervisory control (HSC). These actions are the topic of the present report.

Section one opens the discussion with a review of results from the Patriot Vigilance project. This effort was concerned with the human performance contributors to fratricides involving the Patriot air defense missile system during Operation Iraqi Freedom. Recommendations from the Patriot Vigilance project indicated that the AMD community must address two primary problems associated with automation as applied in current Patriot operations: (1) effective HSC and (2) the level of operator expertise required to employ a highly automated system such as Patriot. These issues are also relevant to AMD systems under development and to other Army systems being fielded to support network-centric warfare concepts. Section two is explicitly concerned with effective HSC. The concept is explored and a definition of effective HSC for AMD operations is provided. Section three examines the impact of automation on human performance. Issues that are addressed in this regard include (1) operator decision-making processes in real-time C2, (2) how smart technology can make operators ineffective as decision makers, and (3) human cognitive vulnerabilities that can add risk to automated C2 operations. Section four then addresses system design and use to facilitate effective HSC. Topics that are discussed in this section include (1) appropriate levels of automation, (2) automation reliability and its impact on performance, and (3) design to support enhanced SA. Finally, Section five proposes a path forward for AMD. Issues that are addressed in this respect include: (1) the practical limits of automation, (2) limits of human expertise, (3) tradeoff and policy considerations, and (4) specific actions to facilitate effective HSC. These actions are: (1) automate only when justified and then carefully, (2) consider adaptive automation when feasible and practical, (3) be brutally honest about automation reliability, (4) provide SA support rather than decisions, (5) use automation for assistance in carrying out routine and low-level actions rather than high-level cognitive tasks, (6) increase the level of crew and battle staff expertise, (7) be aware that there are limits to these potential solution sets—design and human expertise, and (8) resist the temptation to place C2 emphasis on the “gizmo” (C2 technology) rather than on the person using the gizmo. A framework for implementing these actions and proactively managing the human performance risks associated with real-time C2 in AMD operations is also presented and discussed.

1. Background

1.1 Overview

One of the defining properties of the next generation of air and missile defense (AMD) command and control (C2) systems is an increasing reliance on automation. This report is the second in a series of three dealing with human performance and training issues in the development and effective use of automated AMD C2 systems. The first report (Hawley, Mares, & Giammanco, 2005) addressed the impact of automation on air defense operators and the consequences of their role change from traditional operators to supervisory controllers. The present report expands upon that original material and addresses the issue of developing effective human supervisory control (HSC) in AMD C2 systems. Together, these reports are intended as a primer on automation, supervisory control, and effective human performance for commanders, concept developers, system designers, trainers, and other personnel involved with decision making and operations for the next generation of AMD C2 systems.

1.2 Concepts and Terms

Sheridan (1992, p. 3) defines automation as "the automatically controlled operation of an apparatus, a process, or a system, by mechanical or electrical devices that take the place of human organs of observation, decision, or effort." In contemporary AMD C2, the combination of operational complexity and technical advances have led to a situation in which functions—perception, decision-making, response selection and implementation–assigned to the human subsystem in previous generations of AMD systems are now assigned to the machine subsystem.

Supervisory control is defined as a situation in which "one or more operators are continually programming and receiving information from a computer that interconnects through artificial effectors and sensors to the controlled process or task environment" (Sheridan, 2002, p. 115). Under a supervisory control regimen, operators do not interact with the controlled process directly, as they previously did in manual or less automated systems. Rather, the operators receive information from and provide input to a computer which, in turn, directs the controlled process. The operator's role is thus changed from direct, on-line process control to supervisor of a mostly computer-directed process. Their job is to supervise the computer controller. The consequences of this role transformation—though often subtle—must be reflected in system design, performance support feature (i.e., job aiding), and operator-controller[1] training. The challenge facing AMD system developers and users is how to achieve effective HSC of AMD air

[1]To differentiate traditional manual operators from operators performing in a HSC setting, these latter personnel are referred to as "operator-controllers."

battle operations and management (i.e., battle staff operations) without negating the positive effects of technology and mission-essential automation.

1.3 The Patriot Vigilance Project

Personnel from the Army Research Laboratory's Human Research and Engineering Directorate (HRED) began looking into Patriot and AMD system performance at the invitation of the then Ft. Bliss Commander, MG Vane. General Vane was interested in operator vigilance and situation awareness (SA) as they relate to the performance of automated AMD battle command systems. (Note: Endsley (1996) defines SA as the perception of elements in the environment, the comprehension of their meaning, and the projection of their status in the near future.) He was particularly concerned by what he termed a "lack of vigilance" on the part of Patriot operators along with an apparent "lack of cognizance" of what was being presented to them on situation displays with an ensuing "absolute trust in automation." The General's request for human engineering support was prompted by an unacceptably high number of actual or near fratricide incidents by Patriot units during Operation Iraqi Freedom (OIF).

The project staff spent most of the summer and fall of 2004 performing a root cause analysis (RCA) of the OIF fratricide incidents—reading documents, interviewing knowledgeable personnel in the Ft. Bliss area, and observing training and operations. An initial report was delivered to MG Vane in October 2004.

The Patriot Vigilance project was not intended as another exercise in Monday morning quarterbacking. Rather, the intent of HRED's research was to look into the deeper story behind events leading to the OIF fratricides from a human performance perspective. The focus was on actionable solutions—the path forward—rather than a further dissection of the incidents of the past. While studying incidents such as the OIF fratricides does create opportunities for rapid learning and organizational change, hindsight is not foresight. After an incident, investigators have all of the critical information necessary to understand what happened. But that information was not available to participants before the fact. In looking back, investigators tend to oversimplify the situation the actual participants faced. This "hindsight bias" can block an investigator's ability to see and understand the deeper story behind the events in question.

A summary of the RCA results from Patriot Vigilance Phase I is presented in figure 1. The first block in the causal network leading to the OIF fratricides is termed "undisciplined automation," defined as the automation of functions by designers and subsequent implementation by users without due regard for the consequences for human performance (Parasuraman & Riley, 1997). Undisciplined automation tends to define the operators' roles as by-products of the automation. Operators are expected to "take care of" whatever the system cannot handle. However, in the case of Patriot, little explicit attention was paid during design and subsequent testing to determining (1) what these residual functions were, (2) whether operators actually could perform them, (3) how they should be trained, or (4) the impact on the overall system's (hardware plus operators) decision making reliability.

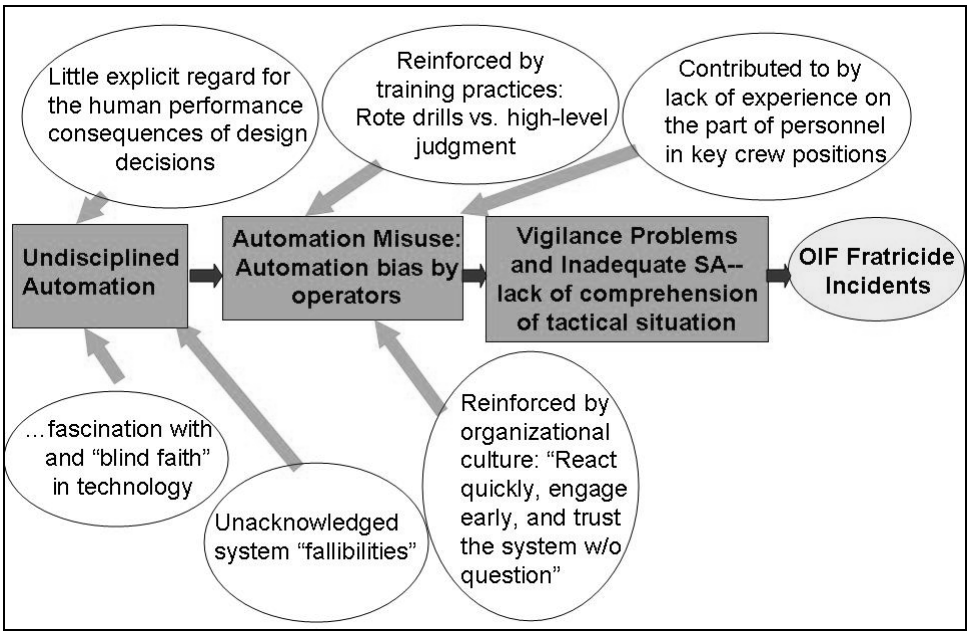

Figure 1. Patriot Vigilance root cause analysis causal network.

The downstream impact of undisciplined automation was exacerbated by two additional factors: (1) unacknowledged system fallibilities, and (2) a fascination with and "blind faith" in technology. A series of Patriot operational tests indicated that the system's engagement logic was subject to track misclassification problems—system fallibilities. However, these sources of automation unreliability were not explicitly patched during system software upgrades, nor did information about them find its way into operator training; battle command practices; tactics, techniques, and procedures (TTPs); or Tactical Standing Operating Procedures (TSOPs). System developers continued to pursue technology-centric solutions to automation reliability problems (increased use of artificial intelligence, non-cooperative target recognition, etc.), but the basic problem remained: The total system (hardware plus crew) was unreliable in critical functional areas, most notably track identification and classification. Users were not informed regarding these problems, or if they were informed, little effective responsive action was taken.

In the aftermath of the first Gulf War (Operation Desert Storm–ODS), the AMD user community acquiesced to the developmental community's apparent lack of concern for problems with Patriot's track classification accuracy. Emboldened by Patriot's seeming success in engaging the Iraqi SCUD threat during ODS, Patriot's organizational culture emphasized "Reacting quickly, engaging early, and trusting the system without question." This cultural norm was exacerbated by the Air Defense Artillery branch's traditional training practices, which have been criticized as emphasizing "rote drills" versus the "exercise of high-level judgment." The user community continued to approach training for Air Battle Operations in much the same manner as March Order and Emplacement or System Set-up. The emphasis was on mastering routines rather than adaptive problem solving. Klein and Pierce (2001) refer to the result of this practice as "experiosclerosis." Crews believe they are experts and "combat ready" because they are good at

the routines, but the routines can prove to be a strait jacket during combat. Traditional individual and unit evaluation practices reinforced this mistaken belief on the part of crews and commanders at all levels by focusing only on satisfactory performance of routine drills.

A second detrimental factor was the Branch's traditional personnel assignment practices which tended to place inexperienced personnel in key crew positions in the Patriot Engagement Control Station (ECS) and Information and Coordination Central (ICC). Before the first round was fired during OIF, the stage was thus set for what Parasuraman and Riley (1997) refer to as "automation misuse," specifically automation bias on the part of Patriot operators. Automation bias is defined as unwarranted over-reliance on automation, and has been demonstrated to result in failures of monitoring (vigilance problems) and accompanying decision biases (an absolute and unthinking trust in automation—let's do what the machine recommends). Recall that these are the very concerns expressed by MG Vane in his kick-off discussion with the Patriot Vigilance staff.

One must be careful, however, not to lay too much blame for these shortcomings at the feet of the Patriot operator-controllers or the battle staff. As suggested in figure 1, the roots of these human shortcomings can be traced back to systemic problems resulting from decisions made years earlier by concept developers, software engineers, procedures developers, trainers, and commanders. In one sense, the OIF Patriot operator-controllers did exactly what they had been trained to do and what Patriot's culture emphasized and reinforced.

Hardware-wise, Patriot is a very lethal system. It can be argued, however, that the system was not properly managed during OIF. Driven by technology and mission expansion, the Patriot crew's role changed from traditional operators to operator-controllers whose primary role is supervision of subordinate automatic control systems. But this role change was not reflected in the AMD culture, design and evaluation practices, battle management concepts, operational procedures, training practices, or personnel usage patterns. Moreover, system management issues (doctrine, battle command concepts, TTPs, TSOPs, etc.) and crewmembers' ability to execute them were not addressed with the same rigor during development and evaluation as hardware and software capabilities. As the lessons of OIF suggest, these aspects of the total "system" are as important to operational effectiveness as hardware and software capabilities.

1.4 Patriot Vigilance Actionable Items and the Defense Science Board Report on Patriot System Performance

HRED's report to MG Vane in October 2004 recommended two primary actionable items to redress the problems discussed in the previous paragraph:

1. Re-define the operators' roles to provide "meaningful human oversight" of system operations, and
2. Develop more effective air battle operations and battle staff personnel—re-look the "level of expertise required to operate such a lethal system on the modern battlefield."

A month later, the Defense Science Board (DSB) (DSB, 2004) echoed HRED's recommendations with the following comments. Although the full DSB report on Patriot system performance is classified, these extracts are not.

> "The Patriot system should migrate to more of a 'man-in-the-loop' philosophy versus a fully automated philosophy—providing operator awareness and control of engagement processes."

and

> "Patriot training and simulations should be upgraded to support this man-in-the-loop protocol including the ability to train on confusing and complex scenarios that contain unbriefed surprises."

The key notion in the first DSB recommendation is captured in the phrase, "providing operator awareness and control of engagement processes." Simply put, soldiers and not the automated system must be the ultimate decision makers in AMD engagements. Decisions to shoot or not to shoot must be made by crews having adequate SA and the expertise to understand the significance of the information available to them. The DSB's first recommendation is synonymous with HRED's first actionable item concerning establishing "meaningful human oversight" of Patriot and other AMD systems.

Prior to proceeding with the current discussion, it is necessary to clarify several terms that are used throughout the report. These terms are "effective HSC," "meaningful human oversight," and "positive human control." The most inclusive term is effective HSC. This means that AMD operator-controllers are able to carry out their explicit role, which is to supervise or direct subordinate automated control systems. Effective HSC implies meaningful human oversight and means that operator-controller supervision of subordinate automated control systems is real and actual and not merely abstract or theoretical. Positive human control is a specific aspect of effective HSC and meaningful human oversight. Positive human control means that engagement decisions—to shoot or not to shoot—are based on conscious problem solving and discernment and not merely the result of automation bias following the system's recommendation.

Putting human decision makers back into the control loop does not mean that the AMD community should try to turn back the clock to the days of Nike Hercules and Hawk and merely re-emphasize traditional manual control strategies and procedures. The situation with Patriot and follow-on AMD systems is too complex for that simplistic solution. Operator-controllers must be augmented by technology in the form of automation. The contemporary AMD environment is simply too complex and demanding to consider any other approach. The DSB report on Patriot system performance voices a similar view:

> "Future conflicts will be more stressing than OIF. They will include ballistic missiles, cruise missiles, UAVs (unmanned aerial vehicles), and enemy aircraft. This will demand a very capable air and missile defense system incorporating robust combat identification and situation awareness."

Automation is a fact of life in AMD operations. The requirement going forward is to approach automation, HSC, and other human systems integration (HSI) issues in a disciplined manner.

The second DSB recommendation having major significance for human performance in contemporary AMD operations concerns operator-controller training and professional development. Here, the DSB was reiterating HRED's conclusion that it is necessary to "re-look the level of expertise necessary to operate such a lethal system on the modern battlefield." In current usage, the term "expertise" refers to a capability for consistently superior performance on a specified set of representative tasks for a domain. Expertise in AMD Air Battle Operations is derived from all aspects of operator-controller job preparation: traditional training (institutional and unit), professional development (self-directed study and professional military education), and relevant on-the-job experience.

The U.S. Navy faced a similar reconsideration of training practices in the aftermath of the shoot-down of the Iranian airbus by the USS Vincennes in 1988. After more than 10 years of research, the Navy reached several conclusions that are also relevant to the contemporary AMD setting. First, the Navy's research indicated that SA is the key factor determining decision quality in battle command. SA is built upon in-depth technical and tactical expertise. The primary implication of this conclusion is that marginally-skilled or apprentice operator-controllers cannot develop the SA necessary for effective supervisory control, regardless of the sophistication of the battle command hardware suite provided to them. Technology is important, but it is only part of the solution. Relevant and in-depth operator expertise is a co-equal factor in developing SA and providing effective human oversight of system operations. Technology can amplify human expertise, but cannot substitute for it.

The Navy also concluded that Aegis operator-controller training must emphasize the development of adaptive decision-making skills. Adaptive decision-making skills, or the ability to think outside the box defined by routine crew drills, are a key aspect of effective operator-controller performance in ambiguous situations. The third major conclusion was that shipboard (i.e., unit) training must address team in addition to individual performance. Competent crews are the basis of effective unit performance, and crews are more than the sum of their individual members.

Finally, the DSB's recommendation to include "unbriefed surprises" in training does not mean that it is sufficient merely to insert anomalous events like those encountered in OIF into training scenarios. In advanced AMD training, the scenario is the curriculum. And to properly prepare operator-controllers for combat, scenario designers must bear in mind that the "surprises" of OIF are representative of a class of potential anomalies. Selected anomalies occurred then; others—some similar, some different–will occur on future battlefields. It is thus necessary that operator-controllers be imbued with a sense of "mindfulness" that automated battle command systems are fallible. The system's recommendations will be correct most but not all of the time. Training must foster the development of the expertise essential to recognize potential anomalies and the

skills necessary to determine an appropriate course of action. Operator-controllers must walk a fine line between blind faith and wholesale mistrust. AMD decision makers must not underestimate the difficulties associated with adequately meeting this training challenge.

Philip Coyle, former Assistant Secretary of Defense and Director of Defense Test and Evaluation, voiced similar concerns in comments on the DSB's report on Patriot system performance (Talbot, 2005):

> "One of the lessons is that the devil is in the details with respect to software. You really have to understand how these computers and software [suites] work…. Since military equipment grows more networked and automated each year, and thus more dependent on software, solving Patriot's problems could be crucial to the future of warfare."

1.5 The Rest of the Report

The concepts for this report and the third grew out of a series of conversations with the Army Training and Doctrine Command (TRADOC) System Manager-Lower Tier (TSM-LT), during which he suggested follow-on reports addressing two additional topics:

- Design for effective HSC of AMD systems
- Training for effective HSC in AMD operations

The present report addresses the first additional topic: Design for effective HSC. In the present context, the term design refers to the development of the hardware subsystem plus associated organizational, battle management, and human support subsystems—the total system package. Note that this definition includes activities performed by both Materiel and Combat Developers, hence our use of the term "developing." The next section begins this discussion by further exploring the concept of effective HSC and formulating a working definition for AMD.

2. Effective Human Supervisory Control in AMD Operations

2.1 Background

The RCA performed as part of the Patriot Vigilance project concluded that vigilance decrements and inadequate SA on the part of Patriot crews were factors in OIF fratricide incidents. Fratricide board of inquiry (BOI) reports also faulted Patriot crews for what was termed "lack of meaningful oversight of system operations." Patriot crews acquiesced to the automated system's recommendations and permitted engagements to occur that, in retrospect, should have been overruled. It is thus reasonable to conclude that the automation and not the human operator-controllers was the ultimate decision maker for these engagements. In this sense, the crews did

not exhibit meaningful human oversight. Put simply, human supervisory control of Patriot operations was, in these instances, not effective.

2.2 The Operator-Controller's Job in an Automated AMD Setting

To bound and define what is meant by the term effective HSC in AMD operations, let us now consider how the introduction of automation changes the operator-controller's job. First, automation is not an all-or-none phenomenon. Rather, contemporary automation is represented by a continuum ranging from total machine control to computer-driven assists that unburden an overloaded operator. Sheridan (1992) presents the taxonomy of potential levels of contemporary automation shown in table 1. Under Sheridan's taxonomy, the current Patriot system implements automation level 5 or 6, depending on the engagement situation. The operator-controller either explicitly authorizes an engagement or implicitly consents to it by not intervening. But, as discussed later in the current section, that variant of HSC is not necessarily effective HSC. The operator-controllers' responsibilities are broader and deeper than simple explicit or implicit consent to a machine recommendation.

Table 1. Levels of contemporary automation.

1. The computer offers no assistance; the operator must do it all.
2. The computer offers a complete set of actions, and…
3. Narrows the selection down to a few, or
4. Suggests one, and
5. Executes that suggestion if the operator approves, or
6. Allows the operator a restricted time to veto before automatic execution, or
7. Executes automatically, then necessarily informs the operator, or
8. Informs the operator only if queried, or
9. Informs the operator after execution if it, the computer, decides to.
10. The computer does everything and acts autonomously, ignoring the operator.

Source: Adapted from Sheridan (1992)

The second introductory point is that all automated systems leave some residual functions for human operator-controllers. Table 2 presents a list of residual human control functions in AMD operations. These AMD-specific functional definitions are taken from Hawley (1994), who adapted them from a more generic set given in Sheridan (1992). AMD operator-controllers perform the residual functions listed in table 2 within the context of the Supervisory Control Operations Cycle shown in figure 2. In an automated control setting, the operator-controller's job is to competently participate in the execution of this process. Taking this line of argument one step further, effective HSC is thus defined in terms of the successful execution of the Supervisory Control Operations Cycle. Successful execution, in present usage, refers to implementing the cycle such that the unit's mission objectives are met.

Table 2. Residual human control functions in air and missile defense.

1. Monitor automatic control: Allocate attention among displays to ensure that the system is operating within performance tolerances.
2. Routine intervention: Intervene to update instructions. Interrupt operations to send commands to the computer when abnormal conditions occur.
3. Select alternative courses of action (COAs): Intervene to change the system's operating mode. Assume direct control of the engagement process when the situation requires such action.
4. Modify control structure or execute new COAs: Instruct the computer on how to change the system's operating mode; update or reset control parameters; modify software; adjust TTP; adapt TSOPs, etc.
5. Plan and initialize: Formulate operating strategies; initialize system. Develop options for achieving system's overall goals.

Source: Hawley (1994) as adapted from Sheridan (1992)

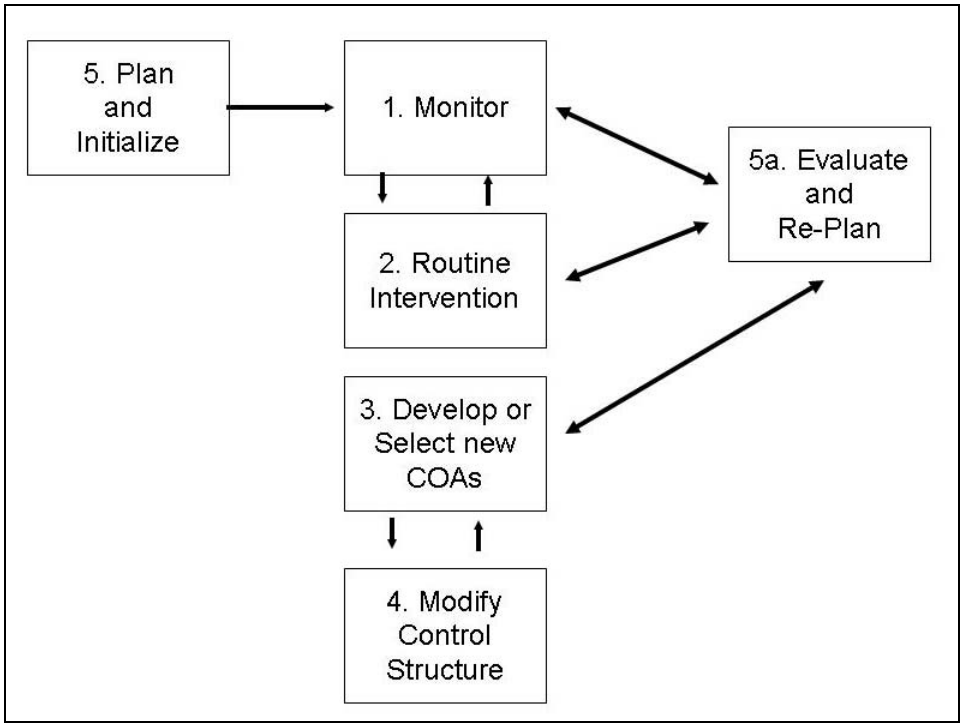

Source: Hawley (1994) as adapted from Sheridan (1992)

Figure 2. Supervisory control operations cycle.

When addressing the "how" of the operator-controllers role in a supervisory control setting, it is instructive to next consider Rasmussen's (1986) supervisory control taxonomy. Under Rasmussen's taxonomy, human tasks in a control system can be classified into one of three categories: denoted skill-based behavior (SBB), rule-based behavior (RBB), and knowledge-based behavior (KBB). Skill-based behaviors consist of sensory and motor performances during acts that, after a statement of intent, take place without conscious control as smooth, automated, and highly integrated behaviors. A simple example of SBB is entering commands into a C2 computer.

In RBB, the task sequence is goal-oriented and consciously controlled by a stored rule: If (Situation)…Then (Action). This governing rule may have been (1) derived empirically during previous operations, (2) communicated from another person's know-how, or (3) prepared on occasion through conscious problem solving and planning. The boundary between SBB and RBB is not distinct; it depends on both the level of training and attention of the operator. Conscious RBB for an inexperienced operator might be automatic SBB for a more experienced one. For example, RBB that is trained to what is termed "automaticity" becomes SBB.

When the operator is faced with situation for which no explicit rules are available, performance control shifts to a higher conceptual level in which actions are goal-oriented and determined on occasion through conscious problem solving and planning. Rasmussen refers to this later category of human performance as KBB. Examples of typical KBB include decision making, planning, and creative thinking. The structure of KBB—that is, how an operator-controller "solves" a problem–is a function of the operator's skill level, experience, and comprehension of the tactical situation.

Rasmussen's supervisory control taxonomy provides a useful perspective on the human performance requirements underlying supervisory versus traditional control. Simply stated, a supervisory control regimen emphasizes and retains operator decision-making and problem-solving tasks while relegating most direct sensory and psychomotor tasks (SBB) and many rule-based performances (RBB) to machine subsystems. Activities in the skill-based performance domain can be allocated either to humans or to the machine. Simple rule-based performances that involve little ambiguity can also be considered for assignment to the machine subsystem.

2.3 A Working Definition of Effective HSC for AMD Operations

Effective HSC requires that engagement decision making be under actual–not theoretical or abstract–human control. Decision to shoot or not to shoot must be made by crews having adequate SA. Crews must have the expertise to understand the significance of the situation display and other information available to them. Additional aspects of effective HSC are that crews must:

1. Be aware of the tactical situation presented on the various situation displays.
2. Have the knowledge, skills, abilities, and experience necessary to translate awareness into understanding.
3. Understand how the technical capabilities and features of the systems they control can impact what is presented to them on situation displays and their ability to engage potential threats.
4. Understand enough of the operational situation (red and blue ground dispositions, red order of battle, etc.) to place the tactical and technical situation in context.

5. Evaluate what their situation displays apparently are telling them against what they should expect to see.

A key notion in the previous list is the idea of context, defined as the interrelated conditions in which something exists or occurs. Woods (2001) comments extensively on the importance of context in human-machine operations. He also discusses the importance of the human operator-controller in providing that context. Woods notes that artificial agents are literal-minded and disconnected from the world, while human agents are context sensitive and have a stake in outcomes. Literal-minded, in present usage, means adhering to facts without nuance. The world is black or white; there is no grey. Literal agents need human help to be grounded in context.

---

3. Automation Issues and Emerging Technologies

3.1 Overview: The Impact of Automation on Human Performance

As discussed in the previous sections, AMD operators' roles have changed from simply "operating" the system against a single air threat—or against multiple threats in a sequential fashion—to managing a resource pool to defeat an enemy air threat. Moreover, the potential air threat is more complex now than in the past, consisting of aircraft, ballistic missiles, cruise missiles, and UAVs. Meeting this emerging threat will require a robust AMD capability—which will bring with it additional automation and automated decision support and will require more effective operator-controllers.

There is a tendency among system developers with little background in human performance issues to assume that automation is innately beneficial. Research in a number of areas (e.g., air traffic control, nuclear power operations, air defense operations, etc.) suggests, however, that such is not always the case. To begin, automation elevates operators into system monitors rather than active controllers. Operators are thus removed from moment-to-moment, active control and become monitors of automated processes. It is a well-established fact that humans make very poor system monitors. Classical vigilance studies have consistently shown that it is difficult even for highly motivated operators to maintain effective visual attention toward a source of information on which little is happening for more than about half an hour (Davies & Parasuraman, 1982).

Beyond classical vigilance, a variety of research has described other problems associated with the role change from traditional to supervisory control. Parasuraman and Riley (1997) argue that automation does not replace human operator tasks. Rather, automation changes the nature of the work that the operators do. It does this in ways that are often unanticipated by system developers and users. Other problems associated with supervisory control generally fall into one of two categories: (1) loss of SA and (2) skill impairment. Automation does not prevent operators from

establishing and maintaining SA, but improper implementation and inadequate training can make it more difficult for them to do so.

3.2 Another View of Operator-Controller Performance in Automated Operations

Operator-controller responses to a control situation are a function of various interacting factors. As shown in figure 3, the primary factors influencing operator-controller responses in a control setting include:

- Immediate demands of the situation and tasks being performed, as perceived by the individual operator-controllers and crews.
- Operator-controller training and experience levels.
- Norms for line operations: "How we do things." The unit's interpretation of TTPs and TSOP guidance.
- Cultural and organizational influences.

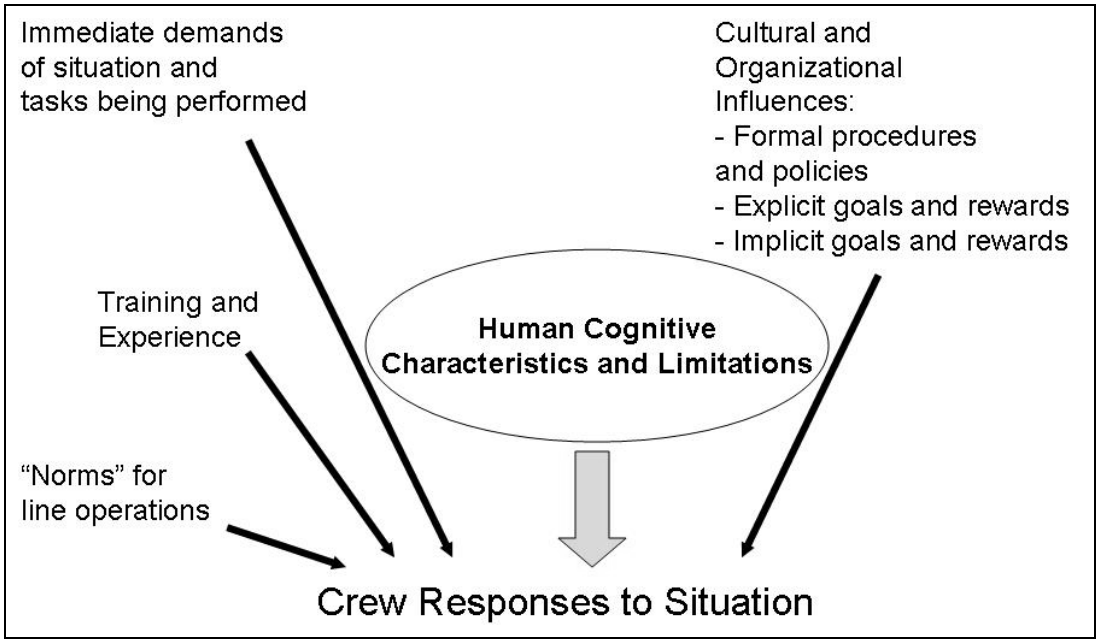

Source: Dismukes and Loukopoulis (2004)

Figure 3. Crew responses are a function of several interacting factors.

The factors depicted in figure 3 interact to shape individual and crew responses to a tactical situation. The term "interact" means that performance is the product of the spontaneous interplay of individual crew members, task, and situation at a particular moment in time. System developers and users also should not underestimate the power of organizational culture and behavioral norms in shaping individual and crew responses and, thus, system performance.

Note in figure 3 that individual and crew responses to a tactical situation also are influenced by what are termed "human cognitive characteristics and limitations." These are addressed in greater depth in the subsections to follow. However, prior to introducing that topic, let us first consider some recent research on how operator-controllers actually make decisions.

3.2.1 How Do Operator-Controllers Make Decisions

The conventional wisdom in most military circles is that decision makers routinely follow what is termed the classical decision making model. That is, the "textbook" way to make an important decision is to (1) list the different options, (2) evaluate those options using a common set of criteria, (3) determine how important each criterion is, (4) rate each option on each criterion, (5) do the math, and (6) determine the optimal choice. The conventional notion of decision making is thorough, systematic, rational, and scientific. But it is also by and large a myth, particularly in real-time C2.

A variety of research and experience indicates that human experts do not follow the classical decision making paradigm in real-time decision making. Rather, human experts use what is termed "pattern matching" to quickly understand a situation and select an appropriate course of action. Klein (2003) refers to this pattern matching-response selection process as Recognition-Primed Decision Making, or RPD. The RPD process works much as described in figure 4. The tactical situation generates a set of cues. This cue set lets the expert recognize a pattern: "I've seen this before." Patterns are associated with what are termed "action scripts": "And these are the actions that I took when I saw this pattern before." Alternative action scripts are assessed using mental simulation based on the expert's mental model of the controlled process. The mental simulation process leads to the selection of one action script as preferred.

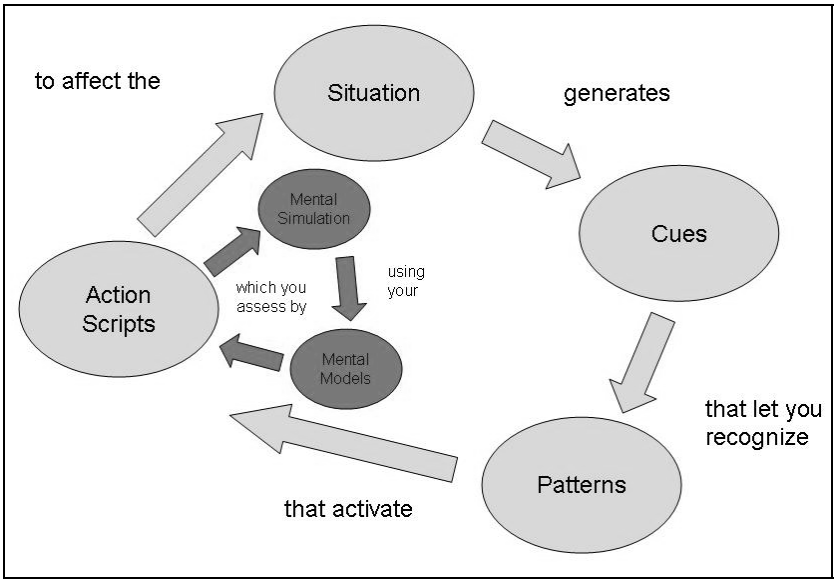

Source: Klein (2003)

Figure 4. The recognition-primed decision model.

A mental model is the expert's internal understanding of how the controlled process "works," or fits together. Norman (2002) remarks that a "good" mental model permits equipment users to predict the effects of their actions. Without a good model, users perform as they are told without really knowing why. As long as things work, they can manage. However, when things go wrong or when the unexpected happens, users frequently are at a loss as to how to proceed.

The RPD cycle is executed rapidly and internally. It is also highly dependent on operator-controller expertise: knowledge (what I know about), skill (what I know how to do), and experience (what I've seen before). Experts immediately "know what to do" in a given situation, and RPD explains how they arrive at a preferred course of action so quickly.

Endsley and her colleagues (Endsley, Bolte, & Jones, 2003) argue that situation awareness, or SA, drives the RPD decision process and is the key factor determining decision quality in battle command. As shown on figure 5, SA is moderated by variables such as preconceptions (what operator-controllers have been led to believe about the system and its capabilities), objectives, ability level, training, experience, doctrine, formal guidance (TTPs and TSOPs), and immediate cognitive workload. Moreover, operator-controller tactical and technical expertise is the dominant factor in establishing Level 2 and Level 3 SA. Comprehension and projection underlie the operator-controller's ability to notice patterns, judge typicality, spot anomalies, and have a "feel" for what is happening around them (Klein, 2003).

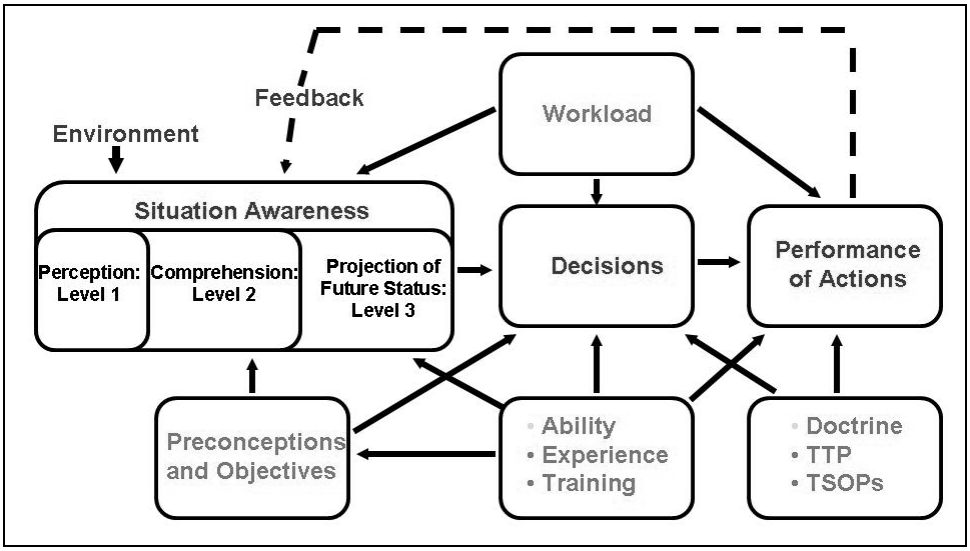

Source: Adapted from Endsley, Bolte, & Jones (2003)

Figure 5. Situation awareness in dynamic decision making.

3.2.2 But "Smart" Technology Can Also Make Us "Stupid"

Earlier in this section, we noted that skill impairment is one of the problems associated with automation. Traditionally, automation research has characterized skill impairment in terms of skill decay–forgetting manual skills that are no longer practiced in an automated setting. The classic example of this type of skill impairment is a pilot who forgets manual control procedures

because of excessive reliance on automated flight control systems. Traditional skill decay or forgetting in an automated setting remains a problem and cannot be discounted if reliance on manual "back-up" skills is an operational necessity.

Recently, however, other researchers such as Klein (2003) have commented on a more subtle and insidious aspect of skill impairment associated with automation. Klein remarks that "smart" technology can make operator-controllers "stupid" in three ways:

1. By disabling the expertise of controllers who are already skilled.
2. By slowing the rate of learning so users do not develop appropriate levels of expertise.
3. By reinforcing dysfunctional skills that will interfere with users' ability to achieve expertise in the future.

In the first case, Klein notes that the information screening and filtering often associated with automation prevents operator-controllers from finding the information they need to make decisions. Operator-controllers thus are denied access to some or all of the decision-cueing information essential to the RPD process. One might say the RPD process can be inadvertently short-circuited through data display and access decisions made by system developers, who often are not familiar with tactical performance requirements.

Second, the same information screening and filtering alluded to in the previous paragraph can deny novice or journeyman controllers access to the data necessary to form the associations between cue sets and environmental patterns. In the first case, information filtering prevents current experts from performing as experts. That same filtering can prevent new operator-controllers from ever developing as experts. They remain intellectually "stunted" creatures of the machine environment.

Klein's third form of skill decay involves reinforcing dysfunctional coping practices within the automated environment. Here, operator-controllers become passive recipients of information from the machine. Klein notes that this passivity tends to make new users reluctant to work around problems or strike out on their own to become true experts. They lack the background to do so. He further remarks that new users lose or never develop their ability to look and search critically within the tactical environment. Smart technology has made them passive and ineffective as problem solvers.

3.3 Varieties of Cognitive Vulnerabilities

Human supervisory controllers are prone to various cognitive vulnerabilities that make it difficult to perform their role as system monitors and supervisors of subordinate automated processes. In this section, we specifically address several of the more commonplace cognitive vulnerabilities in automated operations. Recall that a cognitive vulnerability is something even well-prepared operator-controllers are prone to do behaviorally as a consequence of automation. Failure to take

these tendencies into account during design and later use introduces a potential source of unreliability and risk into system operations.

3.3.1 Automation Bias

Parasuraman and Riley (1997) define automation bias as an unwarranted over-reliance on automation. This bias has been shown to result in failures of monitoring and decision biases, usually in the form of uncritical acquiescence to the automation's recommendations. Automation bias can take the form of both errors of omission (failures to intervene when appropriate) and errors of commission. Omission errors have been shown to be a by-product of vigilance decrements, complacency (a false sense of security), and diffusion of responsibility. In commenting on diffusion of responsibility, Mosier and Skitka (1996) remark that operators frequently exhibit a tendency to off-load their decision-making responsibilities to the machine: "It's probably correct, and who am I to challenge its recommendations"? Organizational influences such as those described earlier for Patriot can reinforce this pass-the-buck tendency on the part of operator-controllers. Sheridan (2002) provides an interesting commentary on how subtly and insidiously automation bias can come about. He notes that repeated successful use leads to trust; trust leads to dependency; and dependency leads to uncritical acquiescence. Operator-controllers are lulled into a false sense of security by previous "successes." Consequently, they do not make checks or ask critical question that in retrospect they "obviously" should have done.

3.3.2 Attentional Tunneling

Endsley, Bolte, and Jones (2003) comment that the stress brought on by surprise, time pressure, and concurrent task demands can be a lethal setup. Stress hampers skilled performance by narrowing operator-controller attention (attentional tunneling) and reducing working memory capacity required to execute even highly practiced tasks. Operator-controllers lose their focus on the big picture and fail to perform activities that they would routinely do under less stressful conditions.

3.3.3 Plan Continuation Bias

Plan continuation bias refers to an unconscious predisposition to continue with an original plan or activity in spite of changing conditions (Dismukes & Loukopoulis, 2004). These authors note that plan continuation bias

1. Is stronger near the completion of an activity.
2. May prevent noticing subtle changes indicating original conditions have changed.
3. May combine with other biases such as frequency sampling (this has always worked before) or cognitive inertia (reactive responding is easier than proactive thinking).

Sheridan poses the question of whether much of contemporary real-time automation illustrates what is referred to as the Roseborough dilemma (as cited in Sheridan, 2002): Automation has been introduced because it can do the job better than a human controller, but the human has been left in the control loop to "monitor" that the automated system is performing correctly and override the automation when it is "wrong." The unstated assumption is that operator-controllers can properly decide when the automation's decisions should be overridden. Humans are expected to compensate for machine unreliability, but they suffer from a variety of cognitive limitations that make it nearly impossible to meet this expectation.

4. Design and Use for Effective Human Supervisory Control

The previous section ended on a cautious note by concluding that while the risks associated with automation unreliability can never be eliminated entirely, they can be managed more effectively. There are a number of positive actions that developers can take to increase overall system reliability. These actions generally fall into one of two categories: (1) level of automation and (2) design for enhanced SA. Each of these topics is discussed in the sub-sections to follow.

4.1 Level of Automation

Hawley, Mares, and Giammanco (2005) note that the dominant human performance theme in automation is function allocation between human operators and machine subsystems. That is, which functions should be automated and to what extent should they be automated. These authors also provide an historical summary of the various approaches to function allocation that have come in and out of vogue since Paul Fitts first formalized the topic in 1951. They conclude by remarking that partitioning control intelligence between human and machine components remains the central theoretical and practical issue in automation and effective HSC.

One of the more promising approaches to function allocation decision making has been advanced by Parasuraman, Sheridan, and Wickens (2000). To begin, these authors define a four-stage model of human information processing, with components as follows:

1. Sensory processing
2. Perception/working memory
3. Decision making
4. Response section

Parasuraman, Sheridan, and Wickens go on to state that the four-stage model of human information processing has its equivalent in system functions that can be automated. These classes of system functions are:

1. Information acquisition: Sensing and registration of input data from the natural environment—what is presented to operator-controllers on the situation display or readily available form other sources.
2. Information analysis: Data analysis and synthesis, along with "cognitive" functions such as inferential processing and projection.
3. Decision and action selection: Selection from among decision alternatives.
4. Action implementation: Execution of the action choice.

For each function class, the level of automated support can vary from "Low" to "High." For decision and action selection, the authors suggest using Sheridan's (1992) Levels of Contemporary Automation (see table 1), where Low is defined as no machine assistance and High as complete machine decision-making autonomy.

Parasuraman, Sheridan, and Wickens suggest that potential levels of automation for various system functions be evaluated in terms of three criteria:

1. The human performance consequences of a decision to automate—the automation's potential impact on:
    a. Operator-controller cognitive workload
    b. Operator-controller complacency, or tendency toward automation bias
    c. Operator-controller skill impairment
2. The automation's reliability in handling that function.
3. The costs of associated decision/action consequences.

To illustrate how Parasuraman, Sheridan, and Wickens' framework for determining level of automation might be applied in AMD, consider an example from Air Traffic Control (ATC). ATC is a performance domain similar in concept to AMD, but with different ends in mind. The results to follow are taken from Wickens (1998) and represent the recommendations of the National Academy of Sciences (NAS) Panel on Human Factors in Air Traffic Control. The Panel was convened to advise the Federal Aviation Administration (FAA) on future directions for ATC automation.

The NAS Panel's recommendations are summarized in table 3. Table 3 makes reference to "conflict resolution" and "maintaining separation." These are the critical functions in an ATC system. Conflicts must be resolved so proper control decisions can be made, and aircraft must be kept separated to avoid mid-air collisions. Collectively, these ATC functions are equivalent to the AMD track identification and classification functions. The capitalized portions in table 3 are included for emphasis.

Table 3. Recommendations for levels of automation in air traffic control.

1. Automation efforts should focus on reliable, high level automation applications for Information Acquisition, Integration, and Presentation and for aiding controller decision making in order to support all system functions. Especially important in the near future is the development of DECISION AIDS for conflict resolution and maintaining separation.
2. The panel recommends implementation of high levels of automation of decision and action selection for system tasks INVOLVING RELATIVELY LITTLE UNCERTAINTY AND RISK. However, for system tasks associated with greater uncertainty and risk, automation of Decision and Action Selection SHOULD NOT PROCEED BEYOND THE LEVEL OF SUGGESTING A PREFERRED DECISION/ACTION ALTERNATIVE. Any consideration for automation above this level must be designed to prevent: LOSS OF VIGILANCE, LOSS OF SA, DEGRADATION OF OPERATIONAL SKILLS, and DEGRADATION OF TEAMWORK AND COMMUNICATIONS. Such design must also ensure the ability to overcome or counteract COMPLACENCY, RECOVER FROM FAILURE, AND PROVIDE A MEANS OF CONFLICT RESOLUTION IF LOSS OF SEPARATION OCCURS.
3. The Panel recommends that the choice of manual (operator initiated) or automatic decision implementation be guided by the level of automation of decision and action selection. Manual implementation is advised at higher levels of automation of decision and action selection, at which automation narrows the decision action alternatives to a few, and more particularly at the level of automation of decision or action selection at which a SINGLE decision/action is selected. MANUAL IMPLEMENTATION WILL ENCOURAGE THE OPERATOR TO REVIEW THE CONTENTS OF THE RECOMMENDED DECISION.
4. The Panel recommends that the availability of computer technology not be reason for automation in and of itself. Clear requirements for functionality that CAN BE ACHIEVED ONLY BY COMPUTER TECHNOLOGY should drive design choices.
5. The Panel recommends that the choice of what functions to automate be guided by recognizing human strengths and the need to compensate for human vulnerabilities.

Source: Wickens (1998)

Three points summarize the Panel's position regarding ATC automation, and apply equally to AMD C2 automation:

1. Do not automate just because you can—because technology seemingly permits it.
2. Automation decisions must explicitly recognize human vulnerabilities and behavioral predispositions.
3. Automation reliability must be a consideration in decisions concerning which system functions to automate and the associated level of automation.

As noted previously, the generalization of these recommendations to AMD are direct. In AMD the function groups requiring careful attention are (1) Information Analysis and (2) Decision and Actions Selection. Under Information Analysis, the specific activities requiring explicit consideration are those associated with Track Evaluation, namely Classification and Identification. Under Decision and Action Selection, critical decision-making activities include those associated with Track Engagement: Engagement Priority and Firing Authorization.

For all of these activities, an analysis following Parasuraman, Sheridan, and Wickens' guidelines suggests low to medium levels of automation. The rationale for this recommendation is as follows: (1) Automation reliability is moderate to high but (2) decision/action consequences also

are high. Automation reliability is noticeably less than 100%, if the events of OIF are an indicator. Moreover, mistakes at any of the stages in the engagement sequence could result in either a fratricide (a false positive–declaring a track hostile when the track actually is friendly) or a missed hostile track (a false negative–declaring a track friendly when the track actually is hostile).

When considering the human performance consequences of recommended levels of automation, high levels of automation for the activities listed above present considerable potential for automation bias as well as longer-term skill impairment. Neither of these possibilities is desirable. On the other hand, low levels of automation for these activities might result in high levels of operator-controller cognitive workload, particularly during intense engagements. In order to maintain acceptable levels of system performance, crews would thus have to be provided with extensive SA support in performing these activities. Adaptive automation, or the ability to change levels of automation in real-time as circumstances require, is yet another performance-enhancing possibility.

This brief example from AMD illustrates that even an objective level of automation analysis might not provide a clear-cut path forward. Tradeoffs and compromises are required. In the previous example, low levels of automation might reduce the potential for automation bias and produce less skill impairment, but result in unacceptable levels of operator workload during intense engagements. Operator-controller workload might be manageable using SA support features or adaptive automation. The operative notion here is "might." Follow-on concept evaluation and operational testing would be required to define and validate an acceptable mix of performance support features. The level of automation analysis does, however, highlight issues requiring attention and suggest possible mitigation strategies.

Given the importance of automation reliability in automation decisions, we next discuss some recent results in that area.

4.2 A Note on Automation Reliability

Reliability is defined as being dependable or capable of being relied upon. Extending this basic definition to real-time C2 as in AMD battle command, a reliable automated system is one in which the functions assigned to the machine are performed accurately and appropriately. Lee and Moray (1992) present results indicating that if operator-controllers have a choice, trust in automation determines usage. Simply put, operator-controllers will elect not to use a system they do not trust. Muir (1988) asserts that trust in automation is affected by the same factors that influence trust between individuals: effectiveness and reliability. In the case of AMD C2, operator-controllers do not have a choice. They must use the system they are provided. So let us next consider the issue of trust in situations where the automation is not perfectly reliable.

Rovira, McGarry, and Parasuraman (2002) remark that one of the true "ironies" of automation is that the more reliable the automation, the greater its detrimental effects when it does fail. As

discussed previously, repeated successful use lulls operator-controllers into a false sense of security or complacency regarding the automation's performance. These authors also report results indicating that operators become over-reliant on automation when it provides decision and action choices for them and do not check underlying information choices as carefully as when only information automation is provided. They conclude that if highly reliable decision automation cannot be guaranteed, then information automation alone should be provided.

Wickens, Dixon, and Ambinder (2005) present results indicating that imperfect automation is *manageable*, but users must be pre-warned of the nature and source of the automation's imperfections. These authors caution, however, that reliabilities less than 75% are worse than no automation at all, and, in fact, can provide users with what they term a "concrete life preserver."

Cohen, Parasuraman, and Freedman (1997) and Masalonis and Parasuraman (1999) argue that trust in automation should not be all-or-none, but graded and differentiated according to the operational context. These authors refer to this as *situation-specific* trust. The automation may work very reliably in certain contexts, in which the operator should use it and trust it. But in certain other cases, that the operator-controller has been trained to look out for, the automation's recommendations may be suspect. Operators should be told to assess the situation and take the action that best suits the context, in their judgment. If operator-controllers can be trained to recognize the appropriate context, then they can know when to trust the automation and when its recommendations should be discounted.

In a similar vein, Lee and See (2004) assert that automation should be designed for *appropriate* as opposed to *greater* trust. These authors go on to state that in situations involving imperfect automation, operator-controller training must emphasize:

1. Expected system reliability
2. The mechanisms underlying potential reliability problems
3. How usage situations interact with the automation's technical characteristics to affect reliability

The lessons of this brief discussion of automation reliability are clear: Developers and users must be brutally honest regarding automation reliability. Extensive tests must be performed to determine those situations in which the automation does not meet design criteria for reliability. Boundaries of successful system performance must be pushed. Moreover, the mechanisms underlying unreliable performance must also be explored. Commanders and operator-controllers must then be apprised of system unreliability patterns and trained in situations that will expose them to system imperfections.

As indicated on figure 1, the Patriot system suffered from various patterns of unreliability—mostly involving track identification and classification—that were known from test results but were not acted upon by the developer or user communities. Operator-controllers were

admonished to "trust the system without question," and training scenarios did not present trainees with situations involving unreliable system performance. Later, during OIF combat operations, it is not surprising that Patriot operator-controllers exhibited the patterns of vigilance decrements, inadequate SA, and over-trust discussed previously, with tragic consequences.

4.3 Design for Enhanced Situation Awareness

The previous discussion has emphasized the importance of SA to effective C2. We are now ready to address the issue of how to enhance SA through design decisions and usage options. Our intent in the discussion to follow is to present an overview of SA design considerations for AMD commanders and decision-makers. In selecting the design principles to highlight, we have included those we have most often seen violated in contemporary AMD systems and those we judge most important to critical design and usage issues going forward. These include topics such as managing situational uncertainty and coping with automation unreliability.

4.3.1 SA Preliminaries

SA researchers frequently remark that design for enhanced SA must begin with a set of preliminary actions. The first of these involves understanding user functions, how they are performed, and why they are performed. For example, Endsley, Bolte, and Jones (2003) devote considerable discussion to the importance of what they term "goal-directed task analysis." The objective of goal-directed task analysis is to understand the user's performance situation—cues, responses, decisions, information sources, and the like.

The information obtained from a goal-directed task analysis must be based on the system's contemporary operating environment and not simply reflect a subjective update of job and task analysis material from a previous era. Users are often tempted to pursue this latter course of action because comprehensive task analysis can be an expensive and time-consuming undertaking. In our view, updating AMD job and task analysis data sources must be a high priority for current systems being used in mission settings different from that originally intended (e.g., Patriot against a tactical ballistic missile [TBM] threat) and for all follow-on systems. Goal-directed task analysis involves much more than simply updating an existing task list.

Second, developers and users must understand system fallibilities and sources of uncertainty in decision making. We noted earlier that acknowledging and understanding system fallibilities and other sources of operational uncertainty is crucial to training commanders and operator-controllers in managing system unreliability. The first two components of SA are perception and comprehension: What is to be displayed? How should that information be displayed? Understanding what makes a situation uncertain is critical to determining display content and granularity to best mitigate that uncertainty.

4.3.2 SA Design Principles

SA design guidelines generally fall into one of five categories, as follows:

1. Physical design
2. Managing situational uncertainty
3. Limiting complexity
4. Using alerts appropriately
5. Supporting team and distributed operations

Specific design guidelines are conceptually derived from the technical definition of SA depicted in figure 5. Readers are referred to that figure as a guide to the following discussion.

The first set of principles concern physical design. The following principles adapted from Endsley, Bolte, and Jones (2003) summarize the wide body of research in design for enhanced SA. Note that these principles focus very much on supporting the three levels of SA–perception, comprehension, and projection: I see it; I understand it; I am able to accurately predict how the situation is going to unfold.

- Organize information around user goals.
- Present Level 2 information directly—support comprehension.
- Provide assistance for Level 3 projections—what is going to happen during the next critical time increment?
- Support global SA—the "big picture." Don't leave operator-controllers or the battle staff with a "key-hole" view of the world.
- Use information filtering carefully—remember that cueing information drives the RPD process. Do not deprive operator-controllers of critical decision cueing elements by inappropriately filtering out or aggregating essential elements of information.

The bottom line with respect to physical design is that system developers must understand the operating environment and the information that operator-controllers and the battle staff must have to perform appropriately. They must craft a system that supports operators in performing their central role.

Second, designers and developers must assist operator-controllers and the battle staff in managing situational uncertainty. Key principles in this category include the following:

- Explicitly identify missing information—make it clear to operator-controllers when critical information elements are missing.
- Support sensor reliability assessment—give operator-controllers explicit information about the reliability of data received from sensors of various kinds.

- Use data salience in support of certainty—provide operator-controllers with some indication of the certainty associated with data displayed to them.
- Represent information timelines—provide operator controllers with some indication of the "age" of information displayed to them. How long has it been since this data entry was received?
- Support assessment of confidence in composite data—when data are aggregated, let operator-controllers know which components are reliable, which should be considered suspect, and how this pattern might impact the reliability of the composite data point.
- Support uncertainty management activities—provide operator-controllers with a practical means (e.g., outside connectivity) to resolve uncertainties and conflicts in the data presented to them.

In many tactical situations, operator-controllers or the battle staff must of necessity integrate incomplete and fragmentary information to make a best judgment. One of the crew's primary roles in system operations is uncertainty management. This involves data reliability assessment and a means for resolving data conflicts. Humans cannot perform the uncertainty management role if all data displayed to them are "equal." They must know what can be trusted and what is suspect.

Third, developers simply must limit display and control complexity. Translating perception into comprehension requires that users assimilate the information presented on the various situation displays. Dense, cluttered, and complex displays can serve to prevent translating perception into understanding. Users must wade through the clutter to decide what is relevant and important. We have observed this phenomenon many times in contemporary AMD system where system and operational complexity result in a proliferation of windows and pull-down menus containing almost countless options. All of this presents even experienced operators with an almost bewildering performance situation.

Several of the elements of limiting complexity include the following:

- Just say "no" to feature creep—creeping featurism results in added complexity, and added complexity leads to skill creep.
- Manage rampant featurism through prioritization and flexibility—know what is important.
- Minimize logic branches.
- Map system functions to user goals and mental models.
- Reduce display density—but do not sacrifice coherence.
- Minimize task complexity—the number and cognitive complexity of actions required to accomplish a task.

Note that a central aspect of managing uncertainty is limiting the amount of information displayed to operator-controllers at any given time. This raises the issue, "How much information is too much"? There is no clear answer to this question. Klein (2003) reports research results involving weather forecasters indicating that, beyond five to ten key pieces of information, additional data did not help and even got in the way and reduced forecasting accuracy. Klein's number (5-10) is not too far removed from Miller's (1956) conclusion that the "magic number" seven plus or minus two serves as a practical limit on human capacity for processing information.

Fourth, developers and users must manage "surprises" by using alerts appropriately. Recall our previous remark that the stress brought on by surprise, time limits, and concurrent task demands can be a lethal setup. In the Patriot Vigilance project, MG Vane asked how to deal with the "23 hours and 59 minutes of boredom followed by one minute of panic" that often characterizes AMD operations. The answer to the General's question centers around managing surprise: Try to prevent that one minute of panic. How can this be done? By providing projection support and appropriate alerts so that surprising situations resulting in panic do not occur. Using alerts appropriately means making them (1) clear and unambiguous, (2) not prone to false alarms, and (3) amenable to rapid assessment and diagnosis.

The final category of means to enhance SA involves team and distributed operations. Here, the research is unequivocal: If team or distributed operations are important to proper system functioning, then it is necessary base team and distributed operations on a common operating picture. Team SA must be developed from a common frame of reference, and geographically separated sub-teams must see the same tactical picture in order to properly coordinate their actions.

5. Unresolved Issues and a Path Forward

Adams (2001) argues that in terms of timelines and complexity, warfare has begun to leave "human space": the traditional four-dimensional battlespace manageable by human senses and cognitive capabilities. He asserts that weapons and other military systems under development will function at increasingly higher levels of complexity and responsibility, and increasingly without meaningful human intervention. Human control will become less direct and more abstract. Abstract, in present usage, means in principle but not in reality.

To illustrate Adams' point, consider the situation with Patriot discussed in section 1. In principle, Patriot operator-controllers were in control of the engagement decision-making process. But were they really in control, or was their control merely theoretical or abstract? One can argue that in those instances positive human control was abstract rather than actual.

Somewhat echoing Adams' view, Klein (2003) remarks that in the evolution of operator responsibilities from traditional operator to manager of multi-mode control operations, SA has taken on increasing importance while simultaneously increasing in difficulty. Klein continues by noting that the emerging operating environment is characterized by increasing cognitive load and complexity. However, operator cognitive capabilities are static and prone to a variety of vulnerabilities and limitations. These constraints mean that operator controllers must of necessity focus on a thinner slice of the total information available.

5.1 The Practical Limits of Automation

Sheridan (2002) comments cautiously on the emerging situation with respect to real-time automation in complex, uncertain environments similar to AMD C2. He notes that increasing operational and system complexity lead to higher levels of uncertainty and ambiguity. These, in turn, create additional opportunities for machine error and less potential for effective human control.

Automation unreliability for a system like Patriot must be addressed on a continuous basis as users gain experience with it. Continuous improvements need not, however, always be handled with "hard" (i.e., hardware or software) changes. Some beneficial changes can be accomplished with "soft" alterations involving procedural modifications or other workarounds. An example from the early history of Patriot will illustrate this point. In the late 1970s and early 1980s when Patriot was first being fielded, concept exploration and evaluation studies performed using the Patriot Tactical Operations Simulator (PTOS) indicated that machine processing of the track identification and classification functions would be problematic (see Hawley, Howard, & Martellaro, 1982). Consequently, a decision was made to designate one of the operator stations in the ECS and ICC as a "friendly protector." The friendly protector's primary role was to manually support the track identification and classification functions for unknown tracks. In essence, this role designation amounted to a *de facto* modification of the Patriot's function allocation scheme that was accomplished using soft versus hard changes.

The friendly protector role designation was de-emphasized in the aftermath of ODS when it was anticipated that Patriot would be used mostly in a counter-TBM role. However, the possibility for similar soft function allocation changes remains an option. This historical note also illustrates the importance of an easily accessible concept evaluation testbed like the PTOS to continuous system improvement and in the evaluation of new systems. Norman (2002) comments that much of good design is evolutionary: The design is tested, problem areas are discovered and modified, and it is continually retested and modified over the course of the system's effective life.

5.2 The limits of Human Expertise

A second potential path forward with respect to maintaining meaningful human control of warfighting systems involves human expertise. We have emphasized several times throughout

the report that operator-controller expertise is critical to establishing and maintaining the SA necessary for effective HSC. However, we also argue that there are practical limits to what can be expected from human operator-controllers. Commenting on combat operations observed during ODS, Cordesman and Wagner (1996) remark that technical advances are used to demand more from operators, and meeting these demands often requires "exceptional human expertise." It thus seems that the greater a system's complexity, the greater the likelihood that designers and users will have to work at or beyond their outer limits of expertise. Norman (2002) refers to this phenomenon as "the paradox of technology": Added functionality generally comes at the price of greater complexity, and complexity results in skill creep.

At the same time that military systems are demanding exceptional expertise of operators, it is necessary to consider the practical limits of training and professional development. First, desirable as it might be, it is simply not possible to train for every potential situation that operator-controllers might face—indeed, many are not known. Moreover, as Kozlowski (1998) and Klein and Pierce (2001) point out, we are uncertain how best to provide generic training and procedures that will work across a broad range of unanticipated situations. Operator-controller expertise is important, but it is not a complete solution to the problem of obtaining reliable human-machine system performance.

5.3 Tradeoffs and Policy Issues

The issues discussed in the previous paragraphs do not paint a satisfying picture of the prospects for effective HSC of AMD operations. Adams (2001, p. 65) supports this conclusion with his remark that, "If the problem is how to maintain meaningful human control of autonomous warfighting systems, no good solution presents itself." The path forward is somewhat uncertain, and likely consists more of tradeoffs and compromises rather than absolutes. Systems can be designed to improve automation reliability and enhance operator-controller SA. Moreover, increasing operator-controller expertise will further increase total system reliability. But the system will not be perfect. The potential for adverse events will continue to exist.

If system developers and commanders accept the proposition that automated AMD C2 systems will never be perfect, it is then reasonable to ask how much system unreliability and its associated decision consequences is tolerable. This question is not intended to be dismissive of fratricide and its often tragic consequences. To put the issue of system unreliability in context, it is instructive to consider historical fratricide rates across system classes. The conventional wisdom in military circles is that fratricide is rare—a nominal two percent rate. And performance expectations are set at this level. However, recent research suggests that fratricide may not be as rare as historically assumed. For example, a study performed by the Congressional Office of Technology Assessment (OTA) in the aftermath of ODS, where 24% of coalition casualties were attributable to fratricide, suggests that 15-20% of total losses may be the historical norm (OTA, 1993). A more recent report prepared for the Joint Staff (Sparta, 2002) puts the figure at 11-16% as a percentage of total casualties. The OTA report further

concludes that the rate observed during ODS may in fact be representative of future conflicts. This increased rate of fratricide was hypothesized to be a function of (1) the increased lethality of precision munitions, and (2) increasing reliance on imperfect sensor data and imperfect classification algorithms in engagement decision making. However, the Sparta report advises caution on the 24% conclusion, remarking that the figure may be biased by a very low overall casualty count. The OTA report concluded that reducing fratricide is desirable and feasible, but eliminating it is not; the Sparta report concurred in this conclusion.

Sheridan (2002) notes that automated systems can be made more reliable by restricting their range of operating circumstances. Restrictive rules of engagement (ROE) might lessen the likelihood of adverse events, but these limitations could take Patriot out of the "out of the fight," so to speak. It is appropriate to ask whether restrictive ROE are the best course of action in all situations. The OTA report cited in the previous paragraph cautions, for example, that overly restrictive ROE may reduce combat effectiveness to such an extent that that casualties inflicted by the enemy may increase more than friendly fire losses are reduced.

5.4 A Path Forward

5.4.1 Overview

Before addressing specific actions in the path forward for AMD, let us first present a perspective on organizational risk management advanced by the British psychologist James Reason (Reason, 1990, 1997). To begin, Reason notes that there are two approaches to risk management, the Person approach and the System approach. The Person approach focuses on errors and procedural violations committed by the people at the "sharp end" of system operations: operator-controllers and the battle staff in the case of AMD. The System approach is based on the notion that people are fallible and errors are to be expected. Hence, errors are viewed as consequences rather than causes of system failure and have their primary origin in "upstream" systemic factors. Error countermeasures are based on the assumption that while we cannot change the human condition, we can change the conditions under which humans work. All high-risk systems possess barriers and safeguards against system failure. When an adverse event occurs, the important issue is not who blundered, but how and why the system's defenses failed.

5.4.2 Reason's Swiss Cheese Model of System Defenses

To further illustrate his approach to error risk management, Reason defines what is termed the "Swiss cheese" model of system defenses. Reason argues that all high-technology, high-risk systems have several layers of defenses against adverse events. The most common of these layers of defenses are (1) the engineered system itself, (2) people, and (3) procedures and administrative controls (e.g., TTPs and TSOPs). In an ideal world, each defense layer would be intact (error-proof). In reality, however, the layers are more like slices of Swiss cheese having many holes that are continually opening, shutting, and shifting their location as a function of the operating situation. The presence of holes in any one "slice" does not normally cause an adverse

outcome. Usually, adverse events happen when the holes in several layers line up momentarily to open a "trajectory" for system failure, as shown in figure 6.

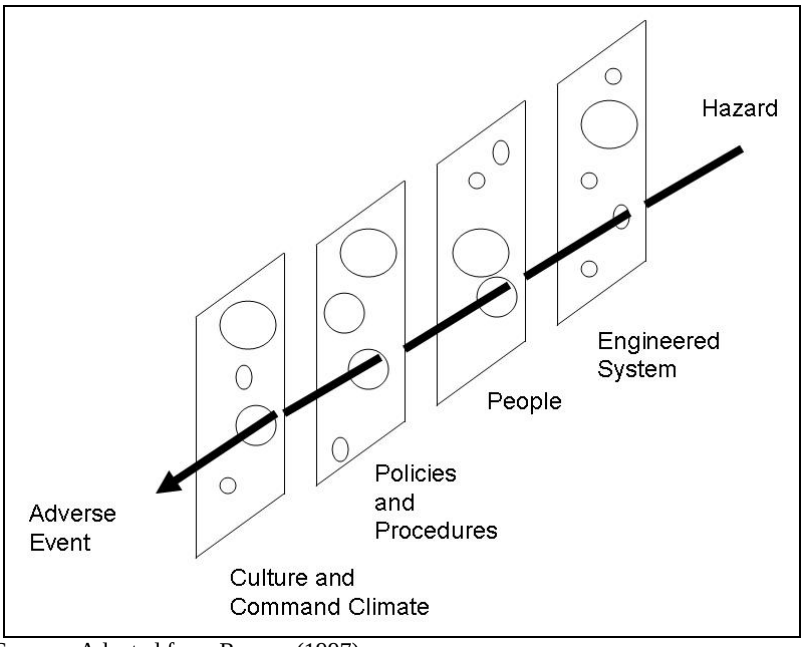

Source: Adapted from Reason (1997)

Figure 6. Reason's Swiss cheese model of organizational defense.

5.4.3 Active and Latent System Failures

Reason continues by noting that the holes in organizational defenses arise for two reasons: active failures and latent conditions. Active failures are actions committed by people on the line. These include "human errors" such as slips, performance lapses, mistakes, procedural violations, and the like. Latent conditions are the inevitable "residual fallibilities" within the system. They arise from decisions made by designers, software engineers, procedures developers, and commanders. Decisions at all of these points have the potential for introducing latent errors into the system. Nearly all adverse outcomes involve a combination of these two sets of factors.

Latent conditions have two kinds of negative impacts. First, they can translate into error-provoking conditions for users—time pressure, stress, fatigue, operator inexperience, and so forth. Second, they can create long-lasting holes or weaknesses in the system's defenses—design deficiencies (system fallibilities), inadequate procedures, training deficiencies, and the like. In the authors' view, the Swiss cheese metaphor accurately characterizes what happened to Patriot during OIF, as described in figure 1.

Reason notes that latent conditions can lie dormant within the system for many years before they combine with active failures (things operator-controllers or the battle staff either do or fail to do) and local triggers (the at-the-moment tactical environment) to create the opportunity for an adverse event. Unlike active failures, whose specific forms are random and thus hard to predict, latent failure conditions can be identified and remedied before an adverse event occurs—through

rigorous testing, being brutally honest about system reliability, etc. Understanding error possibilities leads to proactive rather than reactive risk management.

5.4.4 Elements of the Path Forward

In spite of the preponderance of evidence indicating that developing effective HSC is not a clear-cut proposition, there are a number of actions that system developers, commanders, and decision makers can take to improve the situation in automated AMD C2 operations. These steps are consistent with the proactive risk management concept outlined in the previous subsections.

1. Automate only when justified, and then carefully. Undisciplined or "clumsy" automation can disable operator expertise and set up crews for eventual failure. The framework for determining appropriate levels of automation for system functions discussed in section 3 is a good way to proceed in this regard.

2. Consider adaptive automation when feasible and practical. With adaptive automation, operator-controllers are able to decide function allocation on-line and can select from several options. The function allocation problem is, therefore, not one of allocating functions between humans and the machine once and for all, but dynamic allocation and re-allocation in real time as the process unfolds and requires.

3. Be brutally honest about automation reliability. Take care not to provide users with a "concrete life preserver." Commanders and users must be apprised of potential sources of system unreliability and factor this information into their training and operating plans.

4. Provide SA support rather than decisions. Throughout the report, we have presented caution after caution regarding the negative effects of permitting the automation to make decisions—and then relying upon operators to meaningfully "concur" in these decisions. Research and experience suggests that a preferred course of action is to provide SA support so that operator-controllers rather than the automation are enabled in making key decisions. For example, the FAA has chosen to pursue a more controller-centric approach to ATC automation. Automated support is provided for portions of the controller's job—when such support can be provided reliably, but controllers are still viewed as central to the ATC job.

5. Use automation for assistance in carrying out routine and low-level, rule-based actions rather than performing high-level cognitive tasks. In other words, leave complex rule-based and knowledge-based performances to human operator-controllers.

6. Increase the level of crew and battle staff expertise. Operator-controller expertise is essential to establishing level 2 and Level 3 SA, and these are crucial to effective battle command. Also, clearly address issues associated with knowledge ("know about") versus skill ("know how to do"). In-depth expertise includes both knowing about and knowing how to do. Relevant, on-the-job experience is also a factor in effective decision making.

7. Be aware that there are limits to each of these potential solution sets—design and expertise. In spite of best efforts in both areas, automation will never be perfect, and humans will not always intervene appropriately when the automation fails. Potential trajectories for adverse events will continue to exist. Frankly consider whether is it more useful to adopt reasonable expectations regarding system and operator performance than to persist in unrealistic expectations followed by "surprise" and recriminations when these expectations are not met.
8. Resist placing C2 emphasis on the "gizmo" (the technology) rather than on the person using the gizmo. Wallace (2005, p. 2) cautions that "the network-centric concept introduces a dangerous temptation to shift responsibility for making military decisions from people to the systems themselves."

We are aware that this report raises more questions than it answers. There are, however, advantages to putting problems squarely in the center of the table so that rational and informed discussion can take place. And while we might not have provided satisfactory answers to the questions that gave rise to this report, we have tried to put the underlying issues in perspective. In a recent article in *Harvard Business Review*, Darling, Perry and Moore (2005) discuss the Army's successful use of the After Action Review (AAR) process to facilitate organizational learning. These authors remark that while it is important to correct *things*, it is more important to correct *thinking*. They go on to assert that flawed thinking is the most common cause of flawed execution. Technical corrections affect only the problem that is fixed, but thought-process corrections affect the organization's ability to plan, adapt, and succeed in future actions. In writing this series of reports, one of our objectives is to change the Branch's thinking about automation, the mechanisms underlying effective HSC, and training for effective HSC. After 30 years and two wars in which Patriot was used in combat operations, the time has come to think more realistically about these topics and adapt design, system evaluation, training, performance certification, and usage practices to reflect research results and experience.

6. References

Adams, T. K. Future warfare and the decline of human decisionmaking. *Parameters, Winter* **2001-02**, pp. 57–71.

Cohen, M. S.; Parasuraman, R.; Freeman, J. T. *Trust in decision aids: A model and its training implications*; Technical Report; Cognitive Technologies, Inc: Arlington, VA, 1997.

Cordesman, A. H.; Wagner, A. R. *The lessons of modern war—Vol. 4: The gulf war*. Boulder, CO: Westview Press, 1996.

Darling, M.; Perry, C.; Moore, J. Learning in the thick of it. *Harvard Business Review* **July–August 2005**, pp. 84–92.

Davies, D.; Parasuraman, R. *The psychology of vigilance*. New York: Academic Press, 1982.

Defense Science Board. *Patriot system performance*. (Final Report of the DSB Task Force on Patriot System Performance). Office of the Undersecretary of Defense for Acquisition, Technology, and Logistics: Washington, DC, 2004.

Dismukes, K.; Loukopoulos, L. *The limits of expertise: The misunderstood role of pilot error in airline accidents*. NASA Ames Flight Cognition Laboratory: Moffat Field, CA, 2004.

Endsley, M. R. Automation and situation awareness. In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and applications*. Hillsdale, NJ: Erlbaum, 1996.

Endsley, M. R.; Bolte, B.; Jones, D. G. *Designing for situation awareness: An approach to user-centered design*. New York: Taylor & Francis, 2003.

Hawley, J. K.; Howard, C. W.; Martellaro, A. J. *Optimizing operator performance on advanced training simulators*; Technical Report; U.S. Army Research Institute for the Behavioral and Social Sciences: Alexandria, VA, 1982.

Hawley, J. K. *Theater high altitude area defense (THAAD) system: Training impact analysis*; Technical Report; COLSA Corporation: Huntsville, AL, 1994.

Hawley, J. K.; Mares, A. L.; Giammanco, C. A. *The human side of automation: Lessons for air defense command and control*; ARL TR 3468; Adelphi, MD: U.S. Army Research Laboratory: Adelphi, MD, 2005.

Klein, G. *Intuition at work*. New York: Doubleday, 2003.

Klein, G.; Pierce, L. Adaptive teams. *In Proceedings of the 6th International Command and Control Research and Technology Symposium*, Annapolis, MD, 2001.

Kozlowski, S.W.J. Training and developing adaptive teams: Theory, principles, and research. In J. Cannon-Bowers & E. Salas (Eds.), *Making decisions under stress: Implications for individual and team training*. American Psychological Association: Washington, DC, 1998.

Lee, J. D.; Moray, N. Trust, control strategies, and allocation of function in human-machine systems. *Ergonomics* **1992**, *35*, 1243–1270.

Lee, J. D.; See K. A. Trust in automation: Designing for appropriate reliance. *Human Factors* **2004**, *46* (1), 50–80.

Masalonis, A. J.; Parasuraman, R. *Applying the trust concept to the study of human-automation interaction*. Cognitive Sciences Laboratory, The Catholic University of America: Washington, DC, 1999.

Miller, G. A. The magic number seven plus or minus two: some limits on our capacity for processing information. *Psychological Review* **1956**, *63*, 81–97.

Mosier, K. L.; Skitka, L. J. Automation use and automation bias. *In Proceedings of the Human Factors and Ergonomics Society 40th Annual Meeting* (pp. 204–208). Human Factors and Ergonomics Society: Santa Monica, CA, 1996.

Muir, B. M. Trust between humans and machines and the design of decision aids. In E. Hollnagel, G. Mancini, & D. D. Woods (Eds.), *Cognitive engineering in complex dynamic systems*. London: Academic Press, 1988.

Norman, D. A. *The design of everyday things*. New York: Basic Books, 2002.

Office of Technology Assessment. *Who goes there: Friend or foe?* (OTA-ISC-537). U.S. Government Printing Office: Washington, DC, 1993.

Parasuraman, R.; Riley, V. Humans and automation: Use, misuse, disuse, abuse. *Human Factors* **1997**, *39* (2), 230–252.

Parasuraman, R.; Sheridan, T. B.; Wickens, C. D. A model for types and levels of human interaction with automation. *IEEE Transactions on systems, Man and Cybernetics—Part A: systems and Humans* **2000**, *30* (3), 286–297.

Rasmussen, J. *Information processing and human-machine interaction: An approach to cognitive engineering*. New York: North Holland, 1986.

Reason, J. T. *Human error*. New York: Cambridge University Press, 1990.

Reason, J. T. *Managing the risks of organizational accidents*. Aldershot: Ashgate, 1997.

Rovira, E.; McGarry, K.; Parasuraman, R. Effects of imperfect automation on decision making in command and control. *In Proceedings of the Human Factors and Ergonomics Society 46th Annual Meeting* (pp. 428–432). Human Factors and Ergonomics Society: Santa Monica, CA, 2002.

Sheridan, T. B. *Telerobotics, automation, and human supervisory control*. Cambridge, MA: MIT Press, 1992.

Sheridan, T. B. Humans and automation: System design and research issues. New York: Wiley, 2002.

Sparta, Inc. (2002, March). Combat identification fratricide research study (Final Report). Arlington, VA: Author.

Talbot, D. Preventing "fratricide." Technology Review. Retrieved August 24, 2001, from http://techreview.com/articles/05/06/issue/brief_fratricide.0.asp, 2005, June.

Wallace, W. S. Network-enabled battle command. *Military Review* **May–June 2005**, pp. 2–5.

Wickens, C. D. (Ed.) The future of air traffic control: Human operators and automation. National Academies Press: Washington, DC, 1998.

Wickens, C. D.; Dixon, S. R.; Ambinder, M. S. *Workload and automation reliability in unmanned air vehicles*; Draft Report; U.S. Army Research Laboratory: Adelphi, MD, 2005.

Woods, D. D. Laws that govern cognitive work. Columbus, OH: Cognitive Systems Engineering Laboratory, Ohio State University, 2001.

Distribution List

ADMNSTR
DEFNS TECHL INFO CTR
ATTN DTIC-OCP (ELECTRONIC COPY)
8725 JOHN J KINGMAN RD STE 0944
FT BELVOIR VA 22060-6218

DARPA
ATTN IXO S WELBY
3701 N FAIRFAX DR
ARLINGTON VA 22203-1714

OFC OF THE SECY OF DEFNS
ATTN ODDRE (R&AT)
THE PENTAGON
WASHINGTON DC 20301-3080

US ARMY TRADOC
BATTLE LAB INTEGRATION & TECHL DIRCTRT
ATTN ATCD-B
10 WHISTLER LANE
FT MONROE VA 23651-5850

DIRECTOR
DIRECTORATE OF COMBAT DEVELOPMENTS
ATTN COL H L COHEN
5800 CARTER RD
FT FLISS TX 79916-7001

DIRECTOR
DIRECTORATE OF TRAINING, DOCTRINE, AND LEADER DEVELOPMENT
ATTN COL R BURKE
2 SHERIDAN RD BLDG 2
FT BLISS TX 79916-7001

OFFICE, CHIEF OF AIR DEFNS ARTILLERY
ATTN LTC J H JENKINS III
2 SHERIDAN RD BLDG 2
FT BLISS TX 79916-7001

SMC/GPA
2420 VELA WAY STE 1866
EL SEGUNDO CA 90245-4659

TRADOC SYSTEM MANAGER-LOWER TIER
ATTN COL R JASSEY
BLDG 12, PERSHING RD
FT BLISS TX 79916-7001

TRADOC SYSTEM MANAGER-UPPER TIER
ATTN COL S PETERS
BLDG 12, PERSHING RD
FT BLISS TX 79916-7001

COMMANDING GENERAL
US ARMY AIR DEFNS ARTILLERY CTR AND FT BLISS
ATTN BG R P LENNOX
FT BLISS TX 79916-7001

US ARMY ARDEC
ATTN AMSTA-AR-TD
BLDG 1
PICATINNY ARSENAL NJ 07806-5000

COMMANDING GENERAL
US ARMY AVN & MIS CMND
ATTN AMSAM-RD W C MCCORKLE
REDSTONE ARSENAL AL 35898-5000

US ARMY INFO SYS ENGRG CMND
ATTN AMSEL-IE-TD F JENIA
FT HUACHUCA AZ 85613-5300

US ARMY RSRCH LAB
ATTN AMSRD-ARL-HR-ME A MARES
ROOM 143
FT BLISS TX 79916

US ARMY SIMULATION TRAIN &
INSTRMNTN CMND
ATTN AMSTI-CG M MACEDONIA
12350 RESEARCH PARKWAY
ORLANDO FL 32826-3726

US GOVERNMENT PRINT OFF
DEPOSITORY RECEIVING SECTION
ATTN MAIL STOP IDAD J TATE
732 NORTH CAPITOL ST., NW
WASHINGTON DC 20402

US ARMY RSRCH LAB
ATTN AMSRD-ARL-CI-OK-TP
TECHL LIB T LANDFRIED (2 COPIES)
ABERDEEN PROVING GROUND MD
21005-5066

DIRECTOR
US ARMY RSRCH LAB
ATTN AMSRD-ARL-RO-EN
W D BACH
PO BOX 12211
RESEARCH TRIANGLE PARK NC 27709

US ARMY RSRCH LAB
ATTN AMSRD-ARL-HR-ME
J HAWLEY (10 COPIES)
FT BLISS TX 79916

US ARMY RSRCH LAB
ATTN AMSRD-ARL-D J M MILLER
ATTN AMSRD-ARL-CI-OK-T
TECHL PUB (2 COPIES)
ATTN AMSRD-ARL-CI-OK-TL
TECHL LIB (2 COPIES)
ATTN IMNE-ALC-IMS
MAIL & RECORDS MGMT
ADELPHI MD 20783-1197